Claims 1-17 are pending. Claims 1-4, 7-10, and 13-15 stand rejected under 35 U.S.C. § 102(e) as being anticipated by U.S. Patent No. 6,078,953 to Vaid. Claims 5-6, 11-12, and 16-17 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over U.S. Patent No. 6,078,953 to Vaid in view of U.S. Patent No. 5,276,677 to Ramamurthy.

Reconsideration is requested. No new matter is added. The specification is amended. Claims 1, 7, and 13 are amended. Claims 17-20 are added. The rejections are traversed. Claims 1-20 remain in the case for consideration.

REJECTION OF CLAIMS UNDER 35 U.S.C. § 102(e)

Referring to claim 1, the invention is directed toward a method for enforcing policy over a computer network, the method comprising: defining a template, the template including a set of vectors; assigning a policy to the computer network; monitoring a content stream on the computer network; and enforcing the policy when the content stream is within a threshold distance of the template. Claim 7 is a Beauregard claim paralleling claim 1.

Referring to claim 13, the invention is directed toward an apparatus for enforcing policy over a computer network, the apparatus comprising: a computer; a template stored in the computer, the template including a set of vectors; a policy associated with the template; a monitor installed in the computer adapted to monitor a content stream in the computer network; and a policy enforcer adapted to enforce the policy when the monitor determines the content stream to be within a threshold distance of the template.

In contrast, Vaid teaches a system and method for monitoring quality of service in a network. Vaid uses a traffic management tool coupled to a firewall server. The traffic management tool includes a flow control module and a queueing control module. A bandwidth management tool classifies an information flow into portions, which are directed to the flow control module and the queueing control module.

From the Examiner's citations to Vaid, it appears that the Examiner is misconstruing the term "template" as used in the claims. The term "template" is defined in the specification, at page 3, line 9, as "a set of vectors defined by a semantic content." A sample template is also shown in FIG. 2 of this patent application, and described (at page 4, lines 10-11) as a set of state vectors. Claims 1, 7, and 13 have been amended to describe the template as including a set of vectors. Because this amendment merely clarifies a term already in the claims, these amendments do not narrow the language of the claims: they simply clarify the claim language.

• Vaid has no analog to the concept of a template. The Examiner indicated in marginal notes near column 16, lines 21-22, suggesting that the concept of a "template" could be found in column 16, lines 18-28. But that portion of Vaid only describes classifying the flow of information. For example, "[c]lasses . . . can be defined by source, destination, application, file types, URLs, and other features" (column 16, lines 23-25). Vaid later says, "the flow of information received [is classified] into one of a plurality of predetermined classes" (column 16, lines 27-28).

The point is that in Vaid, classes are predefined, and content is fit into the predefined classes as appropriate. The claimed invention, on the other hand, does not require even one class. Instead, the content is compared mathematically to a set of state vectors (that form a template) to determine a distance between the content and the template. The content is either within a threshold distance of the template, or it is not; when the content is close enough to the template, the policy can be enforced.

The Examiner is also asked to note that the term "distance" appears in independent claims 1, 7, and 13. This term is used in the specification at page 4, lines 17-25. As stated there, distance is measured using a mathematical function. Vaid lacks any description of measuring a distance between the classes and the flow of information. Indeed, it makes no sense to talk about a distance in Vaid, because there is no mathematical definition for either the classes or the flow of information, to which a function could be applied. The best that Vaid can do is to place a flow of information in a particular class and apply the appropriate policy: Vaid cannot measure how well the flow of information fits into the class, or how different two flows of information are assigned to the same class. While the Examiner has made a marginal notation next to column 5, indicating that the Examiner believes "thresholds" are mentioned there, there is no discussion in column 5 (or anywhere else) about distances.

It is worth noting that all of the enabled classes described by Vaid depend on metadata to the content. The described classes (source, destination, file types, URLs, etc.) do not describe the content itself; instead, they relate to external elements. But the invention as claimed describes *monitoring a content stream*. The source, destination, file types, URLs, etc. having nothing to do with the content stream; they relate to the parties and other non-content-specific factors.

The Examiner might be relying on column 13, lines 32-43, where Vaid described an open set of traffic classes. Among these sets is "inbound and outbound information." But even if this description were meant to refer to the actual content of the information flow

(something Vaid does not clearly state), this description is not enabled. Nowhere does Vaid describe how information flows can be grouped into traffic classes based on the actual content in the information flow. Without an enabling description of how to classify flows of information based on the information in the flow, Vaid cannot be said to be teaching classification based on information content. But as this is what is claimed, Vaid cannot anticipate the invention of claims 1, 7, and 13.

As claims 1, 7, and 13 include features not taught or suggested by Vaid, claims 1, 7, and 13 are patentable under 35 U.S.C. § 102(e) over Vaid. Accordingly, claims 1-20 are allowable.

REJECTION OF CLAIMS UNDER 35 U.S.C. § 103(a)

Referring to claim 5, the invention is directed toward a method for enforcing policy over a computer network. Method claim 5 depends from claim 1, and further recites that monitoring a content stream includes: monitoring a portion of the content stream on the computer network; and extrapolating how close the entire content stream is to the template from the portion of the content stream. Claim 11 is a Beauregard claim paralleling claim 5.

Referring to claim 6, the invention is directed toward a method for enforcing policy over a computer network. Method claim 6 depends from claim 1, and further recites that monitoring a content stream includes constructing an impact summary for the content stream. Claim 12 is a Beauregard claim paralleling claim 6.

Referring to claim 16, the invention is directed toward an apparatus for enforcing policy over a computer network. Apparatus claim 16 depends from claim 13, and further recites that the monitor is adapted to monitor only a portion of the content stream on the computer network and extrapolate how close the entire content stream is to the template from the portion of the content stream.

Referring to claim 17, the invention is directed toward an apparatus for enforcing policy over a computer network. Apparatus claim 17 depends from claim 13, and further recites that the monitor includes means for capturing an impact summary for the content stream.

In contrast, Ramamurthy teaches a predictive congestion control for networks. A predictive model is used to predict cross traffic and round trip delay. These data are used to schedule traffic sent to nodes.

To begin, it is important to understand that what Ramamurthy refers to as "cross traffic" has nothing to do with content. Ramamurthy defines "cross traffic" as "real time

traffic, such as voice and video, . . . that cannot tolerate significant delay" (column 3, lines 33-35). Cross traffic reflects types of communication, not content. The same content (e.g., a sound file) could be cross traffic or what Ramamurthy refers to as "controlled traffic": the opposite of cross traffic. For example, if the sound file is part of an on-going communication between two people, it is cross traffic; but if the sound file is a message being retrieved from a voicemail system, it would be controlled traffic. Thus, whether content is determined to be cross traffic or controlled traffic depends on the use of the content, and not on the content itself.

Thus, where Ramamurthy discusses continuous measurement of the cross traffic at nodes in the network (column 11, lines 6-25), Ramamurthy is performing a load analysis, and is not building an impact summary as defined by the claims. Thus, where Ramamurthy talks about "estimat[ing] the total controlled traffic" (column 11, line 23), Ramamurthy is talking about estimating traffic loads, not extrapolating traffic content.

This analysis is pertinent because the Examiner appears to be misinterpreting phrases like "extrapolate how close the entire content stream is to the template" and terms like "impact summary." Based on the citation to column 11, lines 6-25, the Examiner appears to be interpreting extrapolation to refer to guessing how much of the network traffic is controlled traffic versus cross traffic, and to interpreting the term "impact summary" to mean something like "a summary of the impact of the content on the network." Ramamurthy is not applicable because these interpretations are incorrect.

First, an "impact summary" is a set of vectors representing the content in question: in other words, an impact summary has nothing to do with the traffic on the network. The Examiner is referred, for example, to FIG. 3 and the related discussion, which shows measuring a distance between a template and an impact summary. Such a measurement is meaningful only if the impact summary and the template are of the same types, and both the template and the impact summary are shown to be sets of vectors in FIG. 3 and the accompanying description at page 4, lines 26-33. A load analysis cannot be compared, meaningfully or otherwise, to a template because they are of completely unrelated data types, and thus a load analysis of the cross traffic cannot be the same thing as an impact summary.

Second, extrapolation of the content stream refers to extrapolating the content of the entire content stream from the content of the impact summary. If, for whatever reason monitoring the entire content stream is not to be performed, only a portion of the content stream is monitored. The content of the entire content stream is extrapolated from the monitored portion, so that policy enforcement can be performed. The Examiner is referred to
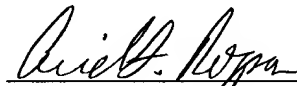
FIG. 5 and the accompanying description at page 6, lines 3-13, which describe exactly this situation.

As claims 5-6, 11-12, and 16-20 include features not taught or suggested by Vaid or Ramamurthy, claims 5-6, 11-12, and 16-20 are patentable under 35 U.S.C. § 103(a) over Vaid in view of Ramamurthy. Accordingly, claims 5-6, 11-12, and 16-20 are allowable.

For the foregoing reasons, reconsideration and allowance of claims 1-20 of the application as amended is solicited. The Examiner is encouraged to telephone the undersigned at (503) 222-3613 if it appears that an interview would be helpful in advancing the case.

Respectfully submitted,

MARGER JOHNSON & McCOLLOM, P.C.

Ariel S. Rogson
Reg. No. 43,054

MARGER JOHNSON & McCOLLOM, P.C.
1030 SW Morrison Street
Portland, OR 97205
503-222-3613
Customer No. 20575